

Network Flow Routing under Strategic Link Disruptions

Mathieu Dahan

Center for Computational Engineering (CCE)
Massachusetts Institute of Technology
Cambridge, MA 02139
Email: mdahan@mit.edu

Saurabh Amin

Department of Civil & Environmental Engineering
Massachusetts Institute of Technology
Cambridge, MA 02139
Email: amins@mit.edu

Abstract—This paper considers a 2-player strategic game for network routing under link disruptions. Player 1 (defender) routes flow through a network to maximize her value of effective flow while facing transportation costs. Player 2 (attacker) simultaneously disrupts one or more links to maximize her value of lost flow but also faces cost of disrupting links. This game is strategically equivalent to a zero-sum game. Linear programming duality and the max-flow min-cut theorem are applied to obtain properties that are satisfied in any mixed Nash equilibrium. In any equilibrium, both players achieve identical payoffs. While the defender’s expected transportation cost decreases in attacker’s marginal value of lost flow, the attacker’s expected cost of attack increases in defender’s marginal value of effective flow. Interestingly, the expected amount of effective flow decreases in both these parameters. These results can be viewed as a generalization of the classical max-flow with minimum transportation cost problem to adversarial environments.

I. INTRODUCTION.

We study a two-player strategic game over a directed network in which player 1 (defender or agency) chooses a flow to be routed from a source node to a destination node, and player 2 (attacker or interdictor) chooses to disrupt one or more edges. The payoff structures in this game are motivated by the previous formulations in both network interdiction problems [1], [2], [3] and attacker-defender games [4], [5], [6]. In our model, we account the value of the effective flow and the transportation cost for the defender, and the value of the lost flow and the cost of attack for the attacker. Specifically, the player 1’s payoff increases in the amount of effective flow that reached the destination node, but decreases with the cost of transporting the initial flow chosen by her. The player 2’s payoff increases in the amount of lost flow as a result of an attack and decreases with the size of attack. This two player game is strictly equivalent to a strictly competitive game (SCG).

We characterize the Nash equilibria for this game for a special class of networks. We show that our game retains some of the nice properties of SCGs. For example, we prove that each player has a unique payoff value in all equilibria. We show that every Nash equilibrium gives a minimaximizing strategy for each player, but the converse is not true. Interestingly, there is a unique maximinimizing strategy but it is not a Nash strategy. Using a combination

of game-theoretic ideas and duality results such as the max-flow min-cut theorem, we are able to compute the values of effective and lost flow and the costs of transportation and attack in terms of the parameters of the game. We also relate the structure of mixed strategy Nash equilibrium of the game to the solutions of the maximum flow minimum transportation cost problem and to the minimum cuts.

Network interdiction problems have already been widely studied, but our focus is to extend these formulations to simultaneous game settings. Related to our approach is the article by Washburn and Wood [7]. The authors model a sequential game, where the defender (leader) chooses one $s-t$ path and then the interdictor (follower) inspects one arc. The objective of this sequential game model is to maximize the probability with which the operator is detected by the interdictor. Our game differs from the model in [7] in that we model a simultaneous game, which captures each player’s strategic uncertainty about her opponent. Secondly, we allow both players to have a much larger set of actions (feasible flow that may contain many $s-t$ paths and loops, and attacks that can disrupt several edges at the same time). Finally, we account for the attacker’s cost of attack as well as the defender’s cost of transporting flow through the network.

Another relevant line of work is by Bertsimas et al. [8]. In this sequential game, the operator first chooses a feasible flow, and then the interdictor disrupts a fixed number of edges. The goal is to minimize the largest amount of flow that reaches the destination node. The authors consider two different models for the disruption: an arc-based formulation where the flow can be rerouted when there is an attack, and a path-based formulation where the flow carried by a disrupted edge is lost. Our formulation is related to the path-based formulation: since we model a simultaneous game, it is reasonable to assume that the flow through disrupted edges is lost and cannot be re-routed. Although, in [8], the interdictor can disrupt several edges at the same time, she must always disrupt the same number of edges for every action, which is still a restriction of the set of actions we considered for our interdictor.

The work of Gueye et al. [4] (also see [9]) is one of the first works to model a simultaneous attacker-defender game,

where the defender (operator) chooses how to send a feasible flow that satisfies the supply and demand constraints (these are exogenous parameters), and the attacker disrupts a single edge of the network. Once again, the utility of the operator only takes into account the amount of flow that reaches the destination node. Similar to [4], our model also considers that the attacker's payoff depends on the amount of lost flow and on the cost of attack. However, in our model, the cost of attack is proportional to the capacities of the edges whereas Gueye et al. only consider a constant cost of attack. Another notable difference is that [4] only considers an uncapacitated graph with given supplies and demands, and we consider a capacitated graph with no constraint on the supplies and demands.

Finally, our model is closely related to the work of Hong and Wooders [10]. The authors model a two player simultaneous game in which an operator chooses a feasible flow and an interdicator disrupts edges of the network, preventing the flow from reaching the destination node. The operator's utility accounts for the amount of flow that reaches the terminal node and the corresponding cost of transportation. However, in [10], the operator's utility is assumed to increase with the cost of attack. This is a central point of departure between their work and ours. Specifically, in our model, the interdicator's utility decreases with the cost of attack. While the results in [10] are based on a max-flow and a min-cut of the graph, we relate Nash equilibria of our game with the solutions of the max-flow with minimum transportation cost problem and with min-cuts.

The rest of the paper is organized as follows: In Section II we discuss the main assumptions and present our game model. Section III presents our main results on the characterization of Nash equilibria of the game. The implications of relaxing some of the modeling assumptions are discussed in Section IV.

II. PROBLEM

A. Preliminaries

Consider a capacitated directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ where \mathcal{V} (resp. \mathcal{E}) represents the set of nodes (resp. the set of edges) of \mathcal{G} . For each edge $(i, j) \in \mathcal{E}$, let $c_{ij} \in \mathbb{R}^+$ denote its capacity, which is the maximum amount of commodity that can pass through (i, j) . Let $s \in \mathcal{V}$ denote a source node and $t \in \mathcal{V}$ a destination node. The flow can only enter the network from s and leave from t , and there is no demand or supply at the nodes that are different from s and t . A flow, denoted x , is a function from \mathcal{E} to \mathbb{R}^+ that assigns to each edge the amount of commodity that goes through it. For notational simplicity, let $x_{ij} := x((i, j))$ denote the flow through edge (i, j) . Let \mathcal{F} denote the set of feasible flows, where a flow x is said to be feasible if it satisfies flow conservation at each node and if the flow through each edge does not exceed its capacity:

$$\begin{aligned} \forall i \in \mathcal{V} \setminus \{s, t\}, \quad \sum_{(j,i) \in \mathcal{E}} x_{ji} &= \sum_{(i,j) \in \mathcal{E}} x_{ij}, \\ \forall (i, j) \in \mathcal{E}, \quad 0 &\leq x_{ij} \leq c_{ij}. \end{aligned}$$

Let Λ denote the set containing all the loops and $s - t$ paths (i.e., a path that starts from s and ends at t) of the network, and x_λ the quantity of flow of x sent through $\lambda \in \Lambda$. Then the edge flows and path/loop flows satisfy:

$$\forall (i, j) \in \mathcal{E}, \quad x_{ij} = \sum_{\{\lambda \in \Lambda \mid (i,j) \in \lambda\}} x_\lambda \quad (1)$$

An $s - t$ cut is a partition $\{S, T\}$ of the set of nodes \mathcal{V} , such that $s \in S$ and $t \in T$. We define the *cut-set* of $\{S, T\}$ as $E(\{S, T\}) = \{(i, j) \in \mathcal{E} \mid i \in S, j \notin S\}$ and the *capacity* of $\{S, T\}$ as $C(\{S, T\}) = \sum_{(i,j) \in E(\{S, T\})} c_{ij}$. Let us recall the *max-flow problem*:

$$\begin{aligned} &\text{maximize} \quad F(x) \\ &\text{subject to} \quad \sum_{(j,i) \in \mathcal{E}} x_{ji} = \sum_{(i,j) \in \mathcal{E}} x_{ij}, \quad \forall i \in \mathcal{V} \setminus \{s, t\} \\ &\quad \quad \quad 0 \leq x_{ij} \leq c_{ij}, \quad \forall (i, j) \in \mathcal{E} \end{aligned}$$

where $F(x) = \sum_{\{i \in \mathcal{V} \mid (i,t) \in \mathcal{E}\}} x_{it}$ denotes the amount of flow passing from the source s to the sink t . The well-known *max-flow min-cut theorem* by Ford and Fulkerson [11] states that the optimal value of the maximum flow problem is equal to the minimum capacity over all $s - t$ cuts. We use Θ to denote the optimal value of the max-flow problem.

We also state the *maximum flow minimum transportation cost problem* by Edmonds and Karp [12] in which the goal is to find a maximum flow with minimum transportation cost:

$$\begin{aligned} (\mathcal{P}) \quad &\text{minimize} \quad \sum_{(i,j) \in \mathcal{E}} b_{ij} x_{ij} \\ &\text{subject to} \quad x \in \mathcal{F} \\ &\quad \quad \quad F(x) \geq F(x'), \quad \forall x' \in \mathcal{F} \end{aligned}$$

where for every edge $(i, j) \in \mathcal{E}$, $b_{ij} \in \mathbb{R}^+$ denotes the cost of transporting one unit of flow through (i, j) . We denote the set of optimal solutions of problem (\mathcal{P}) by Ω .

B. Model

We focus on a simultaneous two-player strategic game $\Gamma := \langle \{1, 2\}, (\mathcal{F}, \mathcal{A}), (u_1, u_2) \rangle$ defined as follows: player 1 (**P1**) is the defender (agency) who chooses to route a flow $x \in \mathcal{F}$ through the network, and player 2 (**P2**) is the attacker (interdicator) who chooses an attack μ to disrupt a subset of edges of graph \mathcal{G} ; see (2). The set of actions for **P1** (resp. **P2**) is given by \mathcal{F} (resp. $\mathcal{A} := \{0, 1\}^\mathcal{E}$).

An attack μ is a function from \mathcal{E} to $\{0, 1\}$ defined as follows:

$$\mu_{ij} := \mu((i, j)) = \begin{cases} 1 & \text{if } (i, j) \text{ is disrupted,} \\ 0 & \text{otherwise.} \end{cases} \quad (2)$$

For our simultaneous game, it is reasonable to assume that after an edge is disrupted (or attacked) by **P2**, the flow that was supposed to cross this edge (if there were no attack) is lost and not re-routed.¹ Thus, the effective flow, denoted x^μ ,

¹We do not consider attacks that can only result in partially disrupted edges and might still permit some flow to pass through the attacked edges.

when the flow x is chosen by **P1** and the attack μ is chosen by **P2** can be expressed as follows:

$$\forall (i, j) \in \mathcal{E}, x_{ij}^\mu = \sum_{\lambda \in \Lambda_{ij}^\mu} x_\lambda,$$

with $\Lambda_{ij}^\mu := \{\lambda \in \Lambda \mid (i, j) \in \lambda \text{ and } \forall (i', j') \in \lambda, \mu_{i'j'} = 0\}$. That is, the effective flow through an edge is the sum of all the initial path flows that go through that edge and that do not contain any attacked edge.

Example 1: Consider the flow chosen by **P1** in Fig. 1a. If **P2** chooses to disrupt edges $(1, t)$ and $(s, 2)$, the unit flows through paths $\{s, 1, t\}$ and $\{s, 2, t\}$ are lost. The resulting effective flow is shown in Fig 1b.




Fig. 1: Example graph.

The payoff of **P1** is defined as follows:

$$u_1(x, \mu) = \underbrace{p_1 F(x^\mu)}_{\text{value of effective flow}} - \underbrace{C_1(x)}_{\text{transportation cost}}, \quad (3)$$

where $p_1 \in \mathbb{R}^+$ is the marginal value of the flow assessed by **P1** and $C_1(x) := \sum_{(i,j) \in \mathcal{E}} b_{ij} x_{ij}$ is the cost of transporting the initial flow x . Thus, when one additional unit of flow reaches t , **P1**'s payoff increases by p_1 and at the same time decreases by its transportation cost.

The payoff of **P2** is defined as follows:

$$u_2(x, \mu) = \underbrace{p_2 F(x - x^\mu)}_{\text{value of lost flow}} - \underbrace{C_2(\mu)}_{\text{cost of attack}}, \quad (4)$$

where $p_2 \in \mathbb{R}^+$ is the marginal value of the lost flow to **P2**, (in general, $p_1 \neq p_2$) and $C_2(\mu) := \sum_{(i,j) \in \mathcal{E}} c_{ij} \mu_{ij}$ is the cost of the attack μ . Thus, if the disruption of an edge induces the loss of one unit of flow, the payoff of **P2** increases by p_2 , and at the same time decreases by the disruption cost, equal to the edge capacity.

We allow both players to randomize over their set of pure actions. Let $\Delta(\mathcal{F})$ and $\Delta(\mathcal{A})$ denote the mixed extensions of **P1** and **P2** pure strategies, respectively, i.e.:

$$\Delta(\mathcal{F}) = \left\{ \sigma^1 \in [0, 1]^{\mathcal{F}} \mid \sum_{x \in \mathcal{F}} \sigma^1(x) = 1 \right\},$$

$$\Delta(\mathcal{A}) = \left\{ \sigma^2 \in [0, 1]^{\mathcal{A}} \mid \sum_{\mu \in \mathcal{A}} \sigma^2(\mu) = 1 \right\}.$$

For notational simplicity, we define $\sigma_x^1 := \sigma^1(x)$ and $\sigma_\mu^2 := \sigma^2(\mu)$. Given any function $\varphi : \mathcal{F} \times \mathcal{A} \rightarrow \mathbb{R}$ and any strategy profile $\sigma = (\sigma^1, \sigma^2) \in \Delta(\mathcal{F}) \times \Delta(\mathcal{A})$, we denote $\mathbb{E}_\sigma[\varphi(x, \mu)] := \sum_{x \in \mathcal{F}} \sigma_x^1 \sum_{\mu \in \mathcal{A}} \sigma_\mu^2 \varphi(x, \mu)$ the expectation

of φ with respect to σ . Then, the respective expected payoffs can be expressed as:

$$U_1(\sigma^1, \sigma^2) = p_1 \mathbb{E}_\sigma[F(x^\mu)] - \mathbb{E}_\sigma[C_1(x)] \quad (5)$$

$$U_2(\sigma^1, \sigma^2) = p_2 \mathbb{E}_\sigma[F(x - x^\mu)] - \mathbb{E}_\sigma[C_2(\mu)] \quad (6)$$

A mixed strategy profile $(\sigma^{1*}, \sigma^{2*}) \in \Delta(\mathcal{F}) \times \Delta(\mathcal{A})$ is a Nash Equilibrium (NE) if:

$$\forall \sigma^1 \in \Delta(\mathcal{F}), U_1(\sigma^{1*}, \sigma^{2*}) \geq U_1(\sigma^1, \sigma^{2*}), \quad (7)$$

$$\forall \sigma^2 \in \Delta(\mathcal{A}), U_2(\sigma^{1*}, \sigma^{2*}) \geq U_2(\sigma^{1*}, \sigma^2). \quad (8)$$

Equivalently, at NE, σ^1 (resp. σ^2) is a best response to σ^2 (resp. σ^1). The support of σ^1 (resp. σ^2) is $\text{supp}(\sigma^1) = \{x \in \mathcal{F} \mid \sigma_x^1 > 0\}$ (resp. $\text{supp}(\sigma^2) = \{\mu \in \mathcal{A} \mid \sigma_\mu^2 > 0\}$).

Our objective is to characterize the set of NE of Γ , denoted Σ_Γ , under the following assumption:

Assumption 1: Let $\alpha := \min_{\lambda \in \Lambda_{\text{path}}} \sum_{(i,j) \in \lambda} b_{ij}$. There exists an optimal solution of (\mathcal{P}) , denoted x^* that takes $s-t$ paths with marginal transportation cost equal to α , i.e.,

$$\exists x^* \in \Omega \mid \forall \lambda \in \Lambda_{\text{path}}, x_\lambda > 0 \implies \sum_{(i,j) \in \lambda} b_{ij} = \alpha,$$

where Λ_{path} is the set containing all the $s-t$ paths of the network.

This assumption implies that if $x^* \in \Omega$ denotes a max-flow with minimum transportation cost, the cost of transporting a unit flow through each $s-t$ path taken by x^* is identically equal to α . In addition, every other path in the network cannot have a smaller marginal transportation cost. Note that the case when every $s-t$ has an identical marginal transportation cost is a special case of this assumption. We illustrate Assumption 1 with the following example:

Example 2: Consider the network flow problem in Fig. 2. There is a unique maximum flow with minimum transportation cost x^* which carries 1 unit of flow through paths $\{s, 2, 4, t\}$, $\{s, 2, 3, t\}$ and $\{s, 1, t\}$. Thus, the total amount of flow is equal to 3 units. In this network, $\alpha = 3$, and each path taken by x^* has a marginal transportation cost equal to 3 so the cost of transporting x^* is equal to 9. The remaining paths that are not taken by x^* are $\{s, 4, t\}$ with a transportation cost equal to 4, and $\{s, 1, 3, t\}$ with a transportation cost equal to 3. Thus, Assumption 1 is satisfied.




Fig. 2: Max-flow with minimum transportation cost (drawn in blue). The labels of each edge correspond to the flow it carries (blue), its capacity (red) and its transportation cost (purple).

Admittedly, on one hand, Assumption 1 restricts the class of networks on which the routing game Γ is played. On the other hand, it permits us to obtain strong results on the equilibrium structure.

Note that Γ can be easily extended to networks with multiple sources and multiple destination nodes. Given such a network, one needs to add an extra source (resp. destination) node and connect it to every existing source (resp. destination) node with an uncapacitated edge of transportation cost equal to 0. This modification gives a new network with single source and single destination. The outcome of the game defined for the original network remains the same as that of the game defined for the new network.

III. CHARACTERIZATION OF NE OF Γ

In this section, we study the properties of the set of NE \mathcal{S}_Γ . First, we argue that Γ is *not* an SCG. Recall that a two-person game is said to be strictly competitive if, when both players change their mixed strategies, either the expected payoffs remain the same or one of the expected payoffs strictly increases and the other strictly decreases. More precisely, Γ can be shown to be not an SCG by arguing that u_1 is not an affine variant of $-u_2$ in the sense of Adler et al. [13]. Although Γ is not an SCG, we can show that Γ is strategically equivalent to a zero-sum game. The strategic equivalence between Γ and a zero-sum game implies that \mathcal{S}_Γ is a convex set.

First, we parametrically solve the game Γ and we present NE whose support involves optimal solutions of (\mathcal{P}) and min-cuts.

Secondly, we focus on the region $p_1 > \alpha$, $p_2 > 1$ which only admits mixed NE and we present our main theorem that provides analytical expressions of certain quantities of interest at any NE. Specifically, we prove that each player has a unique payoff value in all $\sigma^* \in \mathcal{S}_\Gamma$ and we characterize the value of effective (resp. lost) flow and the cost of transportation (resp. cost of attack) in terms of the parameters of Γ : p_1 , p_2 , and α .

Third, we relate the mixed strategy NE of Γ to the solutions of the max-flow minimum transportation cost problem (\mathcal{P}) , and also to the minimum cuts. Finally, we can show that for every NE $(\sigma^{1*}, \sigma^{2*}) \in \mathcal{S}_\Gamma$, σ^{1*} (resp. σ^{2*}) is a minimaximizer for player 1 (resp. player 2); but the converse is not true. Also, for each player there is a unique maximinimizing strategy but it is not a Nash strategy. For the sake of brevity, we have omitted the proofs of several lemmas and propositions.

A. Preliminary results

Let us first present the following lemma which states that **P1** does not choose a flow x that contains a loop.

Lemma 1: *Any flow containing loops is not a best response for **P1**.*

That is, **P1**'s all strategies that contain loops are strictly dominated, and they will not be in the support of a NE. Thus, Λ in (1) can be restricted to the set of $s - t$ paths.

In the rest of this paper, we use the following notations: x^0 the action of not sending flow in the network, x^* an optimal solution of (\mathcal{P}) that satisfies Assumption 1, μ^0 the action of not attacking any edge of the network, and μ^{min} the action that disrupts all the edges of a min-cut of the network.

The following propositions 1–3 provide that, for a given α , the game Γ admits qualitatively different equilibria in regions $0 < p_1 < \alpha$ and $p_2 > 0$ (Region **I**), $p_1 > \alpha$ and $0 < p_2 < 1$ (Region **II**), and $p_1 > \alpha$ and $p_2 > 1$ (Region **III**). These results are summarized by Fig. 3.




Fig. 3: Support of equilibrium strategies in Regions **I-III**.

The following result states that no flow and no attack is the unique NE of Γ in Region **I**:

Proposition 1 (Region I): *If $0 < p_1 < \alpha$, then $\mathcal{S}_\Gamma = \{(x^0, \mu^0)\}$ with $u_1(x^0, \mu^0) = 0$ and $u_2(x^0, \mu^0) = 0$.*

Intuitively, when $0 < p_1 < \alpha$, the marginal value of effective flow that reaches the destination node t is less than the marginal transportation cost for every s - t path. Therefore, **P1** will face negative utility if she sends flow through the network. Thus, in this case, her best response is not to route any flow. Since no flow is sent by **P1**, **P2**'s best response is not to attack, otherwise she faces the cost of attack without any value from lost flow.

Next, for Region **II**, we obtain that max-flow with minimum transportation cost and no attack is a pure NE:

Proposition 2 (Region II): *If $p_1 > \alpha$ and $0 < p_2 < 1$, then $\forall x^* \in \Omega$, $\{x^*, \mu^0\} \in \mathcal{S}_\Gamma$. The corresponding payoffs are $u_1(x^*, \mu^0) = (p_1 - \alpha)\Theta$ and $u_2(x^*, \mu^0) = 0$.*

This result can be explained as follows: on one hand, since **P2**'s valuation of lost flow is small ($p_2 < 1$), for any attack, the utility gained from the lost flow is always lower than the cost of attack. Therefore, **P2**'s best response is not to attack any edge. On the other hand, **P1**'s valuation of effective flow reaching t is higher than the disutility it faces in transportation cost ($p_1 > \alpha$). Since **P2** does not disrupt any edge, every flow sent through the network reaches t ; thus, **P1**'s best response is to send a maximum flow. Among the different maximum flows, the max-flows with minimum transportation cost maximize **P1**'s equilibrium payoff.

Not that if $p_1 = \alpha$ and $p_2 < 1$, then both (x^0, μ^0) and (x^*, μ^0) are NE. The equilibrium payoffs are still $(0, 0)$.

The following proposition 3 focuses on Region **III**:

Proposition 3 (Region III): *If $p_1 > \alpha$ and $p_2 > 1$, then Γ has no pure NE. Furthermore, $\exists \sigma_0 = (\sigma_0^1, \sigma_0^2) \in \mathcal{S}_\Gamma$*

such that $U_1(\sigma_0^1, \sigma_0^2) = U_2(\sigma_0^1, \sigma_0^2) = 0$, and $\text{supp}(\sigma_0^1) = \{x^0, x^*\}$ and $\text{supp}(\sigma_0^2) = \{\mu^0, \mu^{\min}\}$. The corresponding probabilities are given by:

- $\sigma_{x^0}^1 = 1 - \frac{1}{p_2}$, $\sigma_{x^*}^1 = \frac{1}{p_2}$
- $\sigma_{\mu^0}^2 = \frac{\alpha}{p_1}$, $\sigma_{\mu^{\min}}^2 = 1 - \frac{\alpha}{p_1}$.

From this result, we can make several useful observations. First, in contrast to Prop. 1 and Prop. 2, in Region **III**, both players must randomize their actions in any equilibrium. Second, the equilibrium (σ_0^1, σ_0^2) , as defined in Prop. 3, can be obtained from a solution of problem (\mathcal{P}) and a min-cut of the graph \mathcal{G} .² Third, (σ_0^1, σ_0^2) has a particularly simple structure, i.e., **P1** either sends the whole max-flow with minimum transportation cost, or does not send any flow in the network. Similarly, **P2** either disrupts all the edges of a min-cut, or does not attack any edge of the network.

Finally, Prop. 3 is consistent with the game-theoretic intuition: **P1**'s equilibrium strategy σ_0^1 is characterized by the parameter p_2 , and similarly, **P2**'s equilibrium strategy σ_0^2 is characterized by the parameters p_1 and α (which is given and fixed under Assumption 1). This intuition can be further explained as follows: as p_2 increases, $\sigma_{x^*}^1$ decreases while $\sigma_{x^0}^1$ increases. When **P2**'s valuation of lost flow, p_2 , is large, she has more incentive to attack, so any flow sent by **P1** will be more likely to be lost. Thus, **P1** chooses not to send any flow with higher probability than sending x^* . Likewise, as p_1 increases, $\sigma_{\mu^{\min}}^2$ increases while $\sigma_{\mu^0}^2$ decreases. Again, when the marginal valuation of effective flow, p_1 , is large, **P1** will prefer to send as much flow as she can. Thus, **P2** will be more likely to gain by inducing an important loss to **P1** by attacking a min-cut. We illustrate Prop. 1, 2 and 3 with an example.

Example 3: Consider the graph in Fig.4. We can see that $\alpha = 3$, and that the max-flow with minimum transportation cost sends 1 unit of flow through $\{s, 1, 3, t\}$, $\{s, 2, 3, t\}$ and $\{s, 2, 4, t\}$ that induce a transportation cost equal to 3 each. Therefore, Assumption 1 is satisfied. It is easy to see that the min-cut set is given by $\{(1, 3), (2, 3), (2, 4)\}$.




Fig. 4: Example network. Edge capacities and transportation costs are denoted in red and purple labels respectively.

The NE described in Prop. 1, 2 and 3 are illustrated in Fig. 5.

²Recall that, in general, computing an equilibrium and enumerating all equilibria of general two-player non-zero sum games are computationally involved problems [14], [15].




Fig. 5: NE described in Prop. 1, 2 and 3.

B. Main Theorem

We now present our main result which identifies several useful properties satisfied by *all* NE of Γ in Region **III** and under Assumption 1.

Theorem 1: If $p_1 > \alpha$ and $p_2 > 1$, and under Assumption 1, then for any $\sigma^* \in \mathcal{S}_\Gamma$:

- 1) Both players' equilibrium payoffs are equal to 0, i.e.:

$$U_1(\sigma^1, \sigma^2) \equiv 0 \quad (9)$$

$$U_2(\sigma^1, \sigma^2) \equiv 0 \quad (10)$$

- 2) The expected amount of initial flow is given by:

$$\mathbb{E}_{\sigma^*} [F(x)] \equiv \frac{1}{p_2} \Theta \quad (11)$$

and the expected transportation cost is given by:

$$\mathbb{E}_{\sigma^*} [C_1(x)] \equiv \frac{\alpha}{p_2} \Theta \quad (12)$$

- 3) The expected cost of attack is given by:

$$\mathbb{E}_{\sigma^*} [C_2(\mu)] \equiv \left(1 - \frac{\alpha}{p_1}\right) \Theta \quad (13)$$

- 4) The expected amount of effective flow is given by:

$$\mathbb{E}_{\sigma^*} [F(x^\mu)] \equiv \frac{\alpha}{p_1 p_2} \Theta. \quad (14)$$

We refer the reader to Appendix for the proof. From Thm. 1 we observe that although \mathcal{S}_Γ is uncountable (following the convexity of \mathcal{S}_Γ), the equilibrium values of common characteristics, such as expected amount of initial and effective flow, and expected transportation cost to **P1** and the expected cost of attack to **P2**, can be computed in closed-form using the parameters of the game and using the optimal value Θ of the max-flow problem. The payoffs of both players are constant for any NE (ref. (9) and (10)). Below we further explain the implications of Thm. 1.

First, following (11), the expected amount of initial flow is always equal to some fraction of the amount of max-flow, and this expectation decreases with p_2 . Likewise, following (13), the expected cost of attack at any NE is always equal to some fraction of the cost of attacking a min-cut, and this expectation increases with p_1 .

Secondly, following (12), the expected cost of transportation is always equal to some fraction of the cost of transporting x^* (indeed, $C_1(x^*) = \alpha\Theta$ by Assumption 1). Notice that while α always governs the equilibrium strategy of **P2**, it also affects **P1**'s equilibrium transportation cost.

Third, following (14), the expected amount of effective flow is always equal to some fraction of the amount of max-flow, and since $\frac{\alpha}{p_1 p_2} < \frac{1}{p_2}$ this flow is always smaller than the expected amount of initial flow. The expected amount of effective flow decreases when p_1 and/or p_2 increase. This is surprising because one might expect that if the marginal value of effective flow increases, its expectation would increase as well. This result can be explained by noting that when p_1 increases, the disruption caused by **P2** increases, so there is more lost flow and the expected effective flow decreases.

Thm. 1 also enables estimation of the performance metrics such as the expected amount of lost flow and the yield of **P1** in any NE. We define the yield as the ratio of the expected amount of flow that reaches t (effective flow) and the expected amount of flow that is sent through s (initial flow). We have the following corollary:

Corollary 1: *The expected amount of lost flow is given by:*

$$\mathbb{E}_{\sigma^*} [F(x - x^\mu)] \equiv \frac{1}{p_2} \left(1 - \frac{\alpha}{p_1}\right) \Theta, \quad (15)$$

and the expected yield is given by:

$$\frac{\mathbb{E}_{\sigma^*} [F(x^\mu)]}{\mathbb{E}_{\sigma^*} [F(x)]} \equiv \frac{\alpha}{p_1}. \quad (16)$$

From (15), we note that the expected amount of lost flow is always equal to some fraction of the amount of max-flow. The corresponding coefficient increases with p_1 , because when p_1 is large, **P1** sends more flow and **P2** disrupts more edges. However, the coefficient decreases when p_2 increases, mainly because when p_2 is large, **P2** causes more disruption and **P1** sends less flow in the network. Interestingly, we obtain that the expected yield is a constant for all NE, and it only depends on α and p_1 ; neither p_2 nor the maximum amount of flow Θ affect the yield.

Finally, it is easy to check that all of these properties are satisfied by (σ_0^1, σ_0^2) defined in Prop. 3.

C. Additional properties

We now discuss additional properties satisfied by the NE of Γ . The following result relates the NE of Γ with min-cuts of the network \mathcal{G} .

Proposition 4: $\forall (\sigma^{1*}, \sigma^{2*}) \in \mathcal{S}_\Gamma, \forall \mu \in \text{supp}(\sigma^{2*})$:

$$C_2(\mu) \leq C_2(\mu^{\min}) = \Theta,$$

$$\forall (i, j) \in \mathcal{E}, \mu_{ij} = 1 \implies \forall x^* \in \Omega, x_{ij}^* = c_{ij}.$$

From this result, we obtain that, in equilibrium, there will be no attack that requires a cost of attack larger than the cost of attacking a min-cut. In addition, for any NE, an edge is disrupted with positive probability only if it is saturated by every maximum flow with minimum transportation cost. In other words, if there exists at least one maximum flow with minimum transportation cost that does not saturate an edge, then this edge will never be disrupted at equilibrium.

Proposition 5: *Consider a min-cut $E(\{S, T\})$, then:*

$$\forall (\sigma^{1*}, \sigma^{2*}) \in \mathcal{S}_\Gamma, \forall (i, j) \in E(\{S, T\}), \mathbb{E}_{\sigma^*} [x_{ij}] = \frac{c_{ij}}{p_2},$$

Furthermore, for any NE whose support is based on attacks that only disrupt edges of $E(\{S, T\})$ we have:

$$\forall (i, j) \in E(\{S, T\}), \mathbb{P}((i, j) \text{ is disrupted}) = 1 - \frac{\alpha}{p_1}.$$

This proposition tells us that at any NE, the expected amount of flow that goes through any edge of a min-cut is always equal to a fraction of its capacity. In addition, if **P2**'s equilibrium strategy only disrupts edges of one min-cut, then the probability with which an edge is disrupted is constant for all the edges of that min-cut, irrespective of the capacities of these edges.

Corollary 2: $\forall (\sigma^{1*}, \sigma^{2*}) \in \mathcal{S}_\Gamma, \forall \text{ min-cut } E(\{S, T\}), \forall (i, j) \in E(\{S, T\}) : \exists x \in \text{supp}(\sigma^{1*}) \mid x_{ij} > 0,$

This corollary means that for any NE and for any edge of a min-cut, there exists a flow chosen with non-zero probability that passes through it.

All of these properties enable us to derive the following bounds which are tight thanks to Prop. 3:

Proposition 6: *Consider $(\sigma^{1*}, \sigma^{2*}) \in \mathcal{S}_\Gamma$. Then we have the following bounds:*

- If $x^0 \in \text{supp}(\sigma^{1*})$, then $\sigma_{x^0}^{1*} \leq 1 - \frac{1}{p_2}$,
- If $x^* \in \text{supp}(\sigma^{1*})$, then $\sigma_{x^*}^{1*} \leq \frac{1}{p_2}$,
- If $\mu^{\min} \in \text{supp}(\sigma^{2*})$, then $\sigma_{\mu^{\min}}^{2*} \leq 1 - \frac{\alpha}{p_1}$,
- If $\mu^0 \in \text{supp}(\sigma^{2*})$, then $\sigma_{\mu^0}^{2*} \leq \frac{\alpha}{p_1}$.

This proposition gives bounds on the probability with which $x^0, x^*, \mu^0, \mu^{\min}$ are chosen when they are in the support of a NE. From these upper bounds, we see that when p_2 is close to 1, the probability with which x^0 can be chosen becomes very small. On the contrary, it is when p_2 is large that x^* can be chosen with small probability. Similarly, when p_1 is close to α , μ^{\min} can be chosen only with small probability, and when p_1 is large, μ^0 can be chosen only with small probability.

Recall the classical result in SCG that maximinimizers are also NE. In the game Γ , x^0 is a *maximinimizer* for **P1** and μ^0 is a *maximinimizer* for **P2**. However, we already know that if $p_1 > \alpha$ and $p_2 > 1$, there is no pure NE. Hence, the maximinimizing strategies do not give NE, i.e., $(x^0, \mu^0) \notin \mathcal{S}_\Gamma$. Still, we have the following result:

Proposition 7: *Each player's payoffs for both maximinimizing and minimaximizing strategies are equal to the payoff*

at NE, i.e.,

$$\begin{aligned}\max_{\sigma^1} \min_{\sigma^2} U_1(\sigma^1, \sigma^2) &= 0 = \min_{\sigma^2} \max_{\sigma^1} U_1(\sigma^1, \sigma^2) \\ \max_{\sigma^2} \min_{\sigma^1} U_2(\sigma^1, \sigma^2) &= 0 = \min_{\sigma^1} \max_{\sigma^2} U_2(\sigma^1, \sigma^2)\end{aligned}$$

Furthermore, the set of minimaximizers is a superset of \mathcal{S}_Γ , i.e., any NE is a minimaximizer.

IV. RELAXING ASSUMPTION 1

We now discuss an example that does not satisfy Assumption 1. We consider the network given in Fig. 6:




Fig. 6: Network that does not satisfy Assumption 1.

The unique max-flow with minimum transportation cost is drawn in blue in Fig. 6. It sends one unit of flow through path $\{s, 1, t\}$ and one unit through path $\{s, 2, t\}$. They both induce a marginal transportation cost equal to 4. However, $\{s, 1, 2, t\}$ has a marginal transportation cost equal to 3, so Assumption 1 does not hold.

Now, we suppose that $3 < p_1 < 4$ and $p_2 > 1$ then we can prove that $(\sigma^{1*}, \sigma^{2*})$ defined by $\sigma_{x^0}^{1*} = 1 - \frac{1}{p_2}$, $\sigma_{x^1}^{1*} = \frac{1}{p_2}$, $\sigma_{\mu^0}^{2*} = \frac{3}{p_1}$, $\sigma_{\mu^1}^{2*} = 1 - \frac{3}{p_1}$ is a NE (x^1 sends one unit of flow through path $\{s, 1, 2, t\}$ and μ^1 disrupts edge $(1, 2)$).

We can show that, in this case, the NE we presented earlier cannot be found anymore, and the one we found is not supported by any max-flow with minimum transportation cost and any min-cut.

V. APPENDIX: PROOF OF THEOREM 1

We will need the following lemma in the proof of the main theorem:

Lemma 2:

$$\forall (\sigma^{1*}, \sigma^{2*}) \in \mathcal{S}_\Gamma, \mathbb{E}_{\sigma^*} [F((x^*)^\mu)] = \Theta - \mathbb{E}_{\sigma^*} [C_2(\mu)] \quad (17)$$

Proof of Lemma 2: We can find a link between the expected payoffs that we can write in two different ways:

$$\begin{aligned}U_1(\sigma^1, \sigma^2) &= p_1 \mathbb{E}_\sigma [F(x)] - \mathbb{E}_\sigma [C_1(x)] \\ &\quad - \frac{p_1}{p_2} \mathbb{E}_\sigma [C_2(\mu)] - \frac{p_1}{p_2} U_2(\sigma^1, \sigma^2) \quad (18)\end{aligned}$$

$$\begin{aligned}U_2(\sigma^1, \sigma^2) &= -\mathbb{E}_\sigma [C_2(\mu)] + p_2 \mathbb{E}_\sigma [F(x)] \\ &\quad - \frac{p_2}{p_1} \mathbb{E}_\sigma [C_1(x)] - \frac{p_2}{p_1} U_1(\sigma^1, \sigma^2) \quad (19)\end{aligned}$$

Let $\sigma^* = (\sigma^{1*}, \sigma^{2*}) \in \mathcal{S}_\Gamma$. Since $(\sigma_0^1, \sigma_0^2) \in \mathcal{S}_\Gamma$ (Prop. 3), (8) and (6) give us:

$$\begin{aligned}0 &= U_2(\sigma_0^1, \sigma_0^2) \geq U_2(\sigma_0^1, \sigma^{2*}) \\ &= \Theta - \mathbb{E}_{\sigma^*} [F((x^*)^\mu)] - \mathbb{E}_{\sigma^*} [C_2(\mu)]\end{aligned}$$

So we get the first inequality:

$$\mathbb{E}_{\sigma^*} [F((x^*)^\mu)] \geq \Theta - \mathbb{E}_{\sigma^*} [C_2(\mu)] \quad (20)$$

Now, since $(\sigma^{1*}, \sigma^{2*}) \in \mathcal{S}_\Gamma$, (7), (5) and Assumption 1 give:

$$U_1(\sigma^{1*}, \sigma^{2*}) \geq \frac{p_1}{p_2} \mathbb{E}_{\sigma^*} [F((x^*)^\mu)] - \frac{\alpha}{p_2} F(x^*) \quad (21)$$

If we combine (18), (6) and (8), using σ_0^2 , we get:

$$\begin{aligned}U_1(\sigma^{1*}, \sigma^{2*}) &\leq \frac{p_1}{p_2} (C_2(\mu^{min}) - \mathbb{E}_{\sigma^*} [C_2(\mu)]) \\ &\quad - \frac{\alpha}{p_2} C_2(\mu^{min}) \quad (22)\end{aligned}$$

By combining (21) and (22), and using the max-flow min-cut theorem, we obtain:

$$\mathbb{E}_{\sigma^*} [F((x^*)^\mu)] \leq \Theta - \mathbb{E}_{\sigma^*} [C_2(\mu)] \quad (23)$$

Equations (20) and (23) lead to:

$$\mathbb{E}_{\sigma^*} [F((x^*)^\mu)] = \Theta - \mathbb{E}_{\sigma^*} [C_2(\mu)]$$

We are now ready to prove the main result. ■

Proof of Theorem 1: Let $\sigma^* = (\sigma^{1*}, \sigma^{2*}) \in \mathcal{S}_\Gamma$. Let us first show 3), but we will need a few equations before. First, let's prove that $U_1(\sigma^{1*}, \sigma^{2*}) \geq 0$. Equation (7) gives:

$$U_1(\sigma^{1*}, \sigma^{2*}) \geq U_1(x^0, \sigma^{2*}) = 0 \quad (24)$$

If we combine (24) and (22), we obtain:

$$\mathbb{E}_{\sigma^*} [C_2(\mu)] \leq \left(1 - \frac{\alpha}{p_1}\right) C_2(\mu^{min}) \quad (25)$$

In order to get the reverse inequality, let us consider the strategy σ_ϵ^1 defined by $\sigma_{x^*}^{1*} = \frac{1+\epsilon}{p_2}$ and $\sigma_{x^0}^{1*} = 1 - \frac{1+\epsilon}{p_2}$. For an ϵ small enough, (7), (5) and (17) give us:

$$\begin{aligned}U_1(\sigma^{1*}, \sigma^{2*}) &\geq \frac{p_1(1+\epsilon)}{p_2} (F(x^*) - \mathbb{E}_{\sigma^*} [C_2(\mu)]) \\ &\quad - \frac{\alpha(1+\epsilon)}{p_2} F(x^*) \quad (26)\end{aligned}$$

We just need to combine (22) and (26) in order to get:

$$\mathbb{E}_{\sigma^*} [C_2(\mu)] \geq \left(1 - \frac{\alpha}{p_1}\right) F(x^*) \quad (27)$$

Equations (25), (27) and the max-flow min-cut theorem give us 3):

$$\mathbb{E}_{\sigma^*} [C_2(\mu)] = \left(1 - \frac{\alpha}{p_1}\right) C_2(\mu^{min}) = \left(1 - \frac{\alpha}{p_1}\right) \Theta \quad (28)$$

We can now use this equation to prove that **P1**'s payoff is equal to 0 at equilibrium: combining (28) and (22) lead to:

$$U_1(\sigma^{1*}, \sigma^{2*}) \leq 0 \quad (29)$$

Therefore, (29) and (24) give:

$$U_1(\sigma^{1*}, \sigma^{2*}) = 0 \quad (30)$$

Let's now show (11). Similarly, let's first prove that $U_2(\sigma^{1*}, \sigma^{2*}) \geq 0$ using (8):

$$U_2(\sigma^{1*}, \sigma^{2*}) \geq U_2(\sigma^{1*}, \mu^0) = 0 \quad (31)$$

Then, (19), (30) and (28) give:

$$U_2(\sigma^{1*}, \sigma^{2*}) \leq \left(1 - \frac{\alpha}{p_1}\right) (p_2 \mathbb{E}_{\sigma^*} [F(x)] - C_2(\mu^{min})) \quad (32)$$

If we combine (32) and (31), we obtain:

$$\mathbb{E}_{\sigma^*} [F(x)] \geq \frac{1}{p_2} C_2(\mu^{min}) \quad (33)$$

In order to get the reverse inequality, let us consider the strategy σ_ϵ^2 defined by $\sigma_{\mu^0}^2 = \frac{\alpha - \epsilon}{p_1}$ and $\sigma_{\mu^{min}}^2 = 1 - \frac{\alpha - \epsilon}{p_1}$. For an ϵ small enough, (18), (8), (6) and (28) give:

$$U_1(\sigma^{1*}, \sigma^{2*}) \leq \frac{\epsilon}{p_2} C_2(\mu^{min}) - \epsilon \mathbb{E}_{\sigma^*} [F(x)] \quad (34)$$

Equations (34) and (24) give us:

$$\mathbb{E}_{\sigma^*} [F(x)] \leq \frac{1}{p_2} C_2(\mu^{min}) \quad (35)$$

Equations (33), (35) and the max-flow min-cut theorem give us:

$$\mathbb{E}_{\sigma^*} [F(x)] = \frac{1}{p_2} F(x^*) = \frac{1}{p_2} \Theta \quad (36)$$

Likewise, if we combine (36) and (32), we get:

$$U_2(\sigma^{1*}, \sigma^{2*}) \leq 0 \quad (37)$$

Equations (37) and (31) give us:

$$U_2(\sigma^{1*}, \sigma^{1*}) = 0 \quad (38)$$

Thus 1) holds.

Now, if we combine (6), (36), (28) and (38), we show 4):

$$\mathbb{E}_{\sigma^*} [F(x^\mu)] = \frac{\alpha}{p_1 p_2} \Theta \quad (39)$$

Lastly, if we combine (5), (39) and (30), we can show 2):

$$\mathbb{E}_{\sigma^*} [C_1(x)] \frac{\alpha}{p_2} \Theta \quad (40)$$

■

ACKNOWLEDGMENTS

This work was supported in part by FORCES (Foundations Of Resilient CyBer-Physical Systems), which receives support from the National Science Foundation (NSF award numbers CNS-1238959, CNS-1238962, CNS-1239054, CNS-1239166), NSF CAREER award CNS-1453126, and the AFRL LABLET - Science of Secure and Resilient Cyber-Physical Systems (Contract ID: FA8750-14-2-0180, SUB 2784-018400).

REFERENCES

- [1] R. K. Wood, "Deterministic network interdiction," *Mathematical and Computer Modelling*, vol. 17, no. 2, pp. 1 – 18, 1993.
- [2] K. J. Cormican, D. P. Morton, and R. K. Wood, "Stochastic Network Interdiction," *Operations Research*, vol. 46, no. 2, 1998.
- [3] D. Bertsimas, E. Nasrabadi, and S. Stillier, "Robust and adaptive network flows," *Operations Research*, vol. 61, no. 5, pp. 1218–1242, 2013.
- [4] A. Gueye, J. C. Walrand, and V. Anantharam, "Design of network topology in an adversarial environment," in *Decision and Game Theory for Security - First International Conference, GameSec 2010, Berlin, Germany, November 22-23, 2010. Proceedings*, ser. Lecture Notes in Computer Science, T. Alpcan, L. Buttyán, and J. S. Baras, Eds., vol. 6442. Springer, 2010, pp. 1–20.
- [5] W. Szeto, "Routing and scheduling hazardous material shipments: Nash game approach," *Transportmetrica B: Transport Dynamics*, vol. 1, no. 3, pp. 237–260, 2013.
- [6] M. Baykal-Gürsoy, Z. Duan, H. V. Poor, and A. Garnaev, "Infrastructure security games," *European Journal of Operational Research*, vol. 239, no. 2, pp. 469–478, 2014.
- [7] K. W. Alan Washburn, "Two-person zero-sum games for network interdiction," *Operations Research*, vol. 43, no. 2, pp. 243 – 251, 1995.
- [8] D. Bertsimas, E. Nasrabadi, and J. B. Orlin, "On the power of randomization in network interdiction," *CoRR*, vol. abs/1312.3478, 2013. [Online]. Available: <http://arxiv.org/abs/1312.3478>
- [9] A. Laszka and A. Gueye, "Quantifying network topology robustness under budget constraints: General model and computational complexity," in *Decision and Game Theory for Security - 4th International Conference, GameSec 2013, Fort Worth, TX, USA, November 11-12, 2013. Proceedings*, ser. Lecture Notes in Computer Science, S. K. Das, C. Nita-Rotaru, and M. Kantarcioglu, Eds., vol. 8252. Springer, 2013, pp. 154–174.
- [10] S. Hong and M. Wooders, "Strategic Network Interdiction," Vanderbilt University Department of Economics, Vanderbilt University Department of Economics Working Papers 1010, Jun. 2010. [Online]. Available: <http://ideas.repec.org/p/van/wpaper/1010.html>
- [11] L. R. Ford and D. R. Fulkerson, "Maximal flow through a network," *Canadian Journal of Mathematics*, vol. 8, pp. 399–404, 1956.
- [12] J. Edmonds and R. M. Karp, "Theoretical improvements in algorithmic efficiency for network flow problems," *J. ACM*, vol. 19, no. 2, pp. 248–264, Apr. 1972. [Online]. Available: <http://doi.acm.org/10.1145/321694.321699>
- [13] I. Adler, C. Daskalakis, and C. H. Papadimitriou, "A note on strictly competitive games," in *Internet and Network Economics, 5th International Workshop, WINE 2009, Rome, Italy, December 14-18, 2009. Proceedings*, ser. Lecture Notes in Computer Science, S. Leonardi, Ed., vol. 5929. Springer, 2009, pp. 471–474.
- [14] D. Avis, G. D. Rosenberg, R. Savani, and B. von Stengel, "Enumeration of nash equilibria for two-player games," *Economic Theory*, vol. 42, no. 1, pp. 9–37, 2010.
- [15] C. Daskalakis, P. W. Goldberg, and C. H. Papadimitriou, "The complexity of computing a nash equilibrium," *Commun. ACM*, vol. 52, no. 2, pp. 89–97, 2009.